

Applying polynomial filtering to mass preconditioned Hybrid Monte Carlo

Taylor Haar^{a,*}, Waseem Kamleh^a, James Zanotti^a, Yoshifumi Nakamura^b

^a*CSSM, Department of Physics, The University of Adelaide, Adelaide, SA, Australia 5005*

^b*RIKEN Advanced Institute for Computational Science, Kobe, Hyogo 650-0047, Japan*

Abstract

The use of mass preconditioning or Hasenbusch filtering in modern Hybrid Monte Carlo simulations is common. At light quark masses, multiple filters (three or more) are typically used to reduce the cost of generating dynamical gauge fields; however, the task of tuning a large number of Hasenbusch mass terms is non-trivial. The use of short polynomial approximations to the inverse has been shown to provide an effective UV filter for HMC simulations. In this work we investigate the application of polynomial filtering to the mass preconditioned Hybrid Monte Carlo algorithm as a means of introducing many time scales into the molecular dynamics integration with a simplified parameter tuning process. A generalized multi-scale integration scheme that permits arbitrary step-sizes and can be applied to Omelyan-style integrators is also introduced. We find that polynomial-filtered mass-preconditioning (PF-MP) performs as well as or better than standard mass preconditioning, with significantly less fine tuning required.

Keywords: 11.38.Gc, Hybrid Monte Carlo algorithm, Multiple time scale integration

1. Introduction

The steady advance in computing power and algorithmic techniques has enabled lattice QCD simulations to be performed at physical quark masses. Generating configurations at or near the physical point provides a significant computational challenge, and the corresponding need for larger lattice volumes means that these simulations require the use of Petascale computing facilities. Furthermore, the complexity of the algorithms used to generate dynamical gauge fields has also increased, with a corresponding increase in the effort required to tune the associated parameters. This motivates investigations into improving algorithmic efficiency and streamlining the tuning process for lattice QCD configuration generation.

The algorithm of choice for generating gauge fields with dynamical quarks is Hybrid Monte Carlo (HMC) [1]. However, vanilla HMC suffers from critical slowing down: moving to smaller quark masses m results in a dramatic increase in the condition number of the Dirac matrix. This leads to a corresponding increase in computational cost — fits to the cost against quark mass m [2] suggest a m^{-3} dependence. This critical slowing down makes physical point simulations with the vanilla HMC algorithm infeasible.

This has led to the development of a large variety of algorithmic improvements for HMC which are actively used in simulations. The performance at lighter quark masses is improved by techniques such as Hasenbusch mass preconditioning [3], polynomial filtering [4], domain decomposition [5], and rational HMC [6]. These improvement techniques modify the fermion action to either decrease the condition number

*Corresponding author

Email address: taylor.haar@adelaide.edu.au (Taylor Haar)

of the fermion matrix or increase the stability of the inversion, such that coarser step-sizes can be used and hence the computational cost can be reduced. Such techniques can be further improved by modifying the HMC integration scheme, namely using multiple time-scales [7] and higher-order integrators [8]. In particular, the use of one or more filtering terms to perform frequency-splitting and to break up the fermionic determinant into multiple time-scales is critical for light quark mass simulations.

Simulations that include two degenerate quark flavours with a mass approaching or at the physical value typically make use of a hierarchy of two, three or more mass preconditioning terms [9–12]. However, it is challenging to simultaneously tune a large number of Hasenbusch mass parameters. In contrast, the parameter tuning required for multiple polynomial filters is relatively simple [4]. This work investigates the effects of combining polynomial filtering with mass preconditioning (Hasenbusch filtering), in an effort to obtain similar or better performance with a simpler tuning process when compared to plain Hasenbusch filtering. As is the norm [4, 9–15], each action term is placed on a different time-scale according to its respective force in order to minimize the overall cost. We use the number of fermion matrix-vector multiplications as a platform-independent benchmark for comparison.

This paper has two main sections. In section 2, we describe the Hasenbusch and polynomial filtering techniques, how the application of multiple time-scales leads to improved performance, and a procedure for tuning the large number of resultant parameters. Section 3 gives an analysis of the performance of Hasenbusch filtering, polynomial filtering, and polynomial-filtered mass-preconditioning.

The appendices outline some of the more technical details. Appendix A describes the choice of polynomial used for the polynomial filtering. In Appendix B, we construct a generalized multiple time-scale integration scheme that permits an arbitrary choice of step-size, and show that it is area preserving and time reversible as required. Finally, Appendix C derives the molecular dynamics force terms for the two filtering methods under consideration.

2. Method

2.1. HMC

The method of choice for including dynamical fermions in lattice QCD simulations is HMC [1], where successive gauge configurations U are generated by introducing a fictitious conjugate momentum field P then preserving the Hamiltonian

$$H[P, U] = \sum \text{tr}[P^2] + S[U]$$

via Hamilton’s equations, giving integration steps

$$\begin{aligned}\hat{T}[\epsilon] : (P, U) &\rightarrow (P, e^{i\epsilon P} U), \\ \hat{S}[\epsilon] : (P, U) &\rightarrow (P - \epsilon F, U).\end{aligned}$$

Here, $S = S_G + S_F$ is the Euclidean action, ϵ is a step-size and $F = \frac{\partial S}{\partial U}$ is the force term. We use a sequence of these steps, typically of unit length in simulation time, to evolve the system from the state (P, U) to (P', U') ; this is known as a molecular dynamics trajectory. The resultant gauge configuration U' then undergoes a Metropolis acceptance step, with acceptance probability

$$P_{\text{acc}} = \min[1, \exp(H[P, U] - H[P', U'])]. \quad (2.1)$$

The main computational cost for HMC is in calculating the force term $F = \frac{\partial S}{\partial U}$. If we consider the basic 2-flavour fermion action with pseudo-fermion field ϕ and Dirac matrix M ,

$$S_F = \phi^\dagger (M^\dagger M)^{-1} \phi \equiv \phi^\dagger K^{-1} \phi, \quad (2.2)$$

the fermionic force term takes the form

$$F = \frac{\partial S_F}{\partial U} = -\phi^\dagger K^{-1} \frac{\partial K}{\partial U} K^{-1} \phi. \quad (2.3)$$

The costly operation here is calculating $K^{-1}\phi$, i.e. solving $\phi = K\chi$ for χ . Due to the size of K , we invert using iterative sparse matrix techniques such as conjugate gradient or other Krylov-space methods. As we go to smaller quark masses, the condition number of K increases and so the inversion requires more matrix multiplications. At the same time, the size and variance of the force term F increases, which then requires a reduction in the integration step size to maintain a reasonable acceptance rate; this increases the frequency with which the expensive matrix inverse evaluations must be performed. For these reasons, filtering techniques that reduce the frequency of costly matrix inversions are essential at light quark masses.

2.2. Filtering methods

By noting that

$$\det K = \frac{\det[LK]}{\det L} \quad (2.4)$$

for any matrices K, L invertible, we can separate our fermion action into multiple terms,

$$S_{\text{filtered}} = \phi_1^\dagger L \phi_1 + \phi_2^\dagger [LK]^{-1} \phi_2. \quad (2.5)$$

Since L acts as a filter for the fermion matrix K , we call methods that use (2.5) filtering methods. The aim of filtering is to reformulate the fermion action in such a way that the partitioned terms form an approximation to the determinant that is easier to calculate (e.g. by reducing the stochastic noise). Typically, the success of a filtering method in reducing the computational cost of a simulation requires that the force F_1 associated with the ‘filter term’ $\phi_1^\dagger L \phi_1$ is relatively cheap to evaluate, and that the filter provides a reduction in the size of force F_2 for the expensive ‘correction term’ $\phi_2^\dagger [LK]^{-1} \phi_2$.

Mass preconditioning [3] (also known as Hasenbusch preconditioning) is the predominant filtering method used in modern lattice simulations for two degenerate quark flavours, and takes the form

$$S_{\text{MP}} = \phi_1^\dagger J^{-1} \phi_1 + \phi_2^\dagger J K^{-1} \phi_2, \quad (2.6)$$

where $J = W^\dagger W$ and W is a fermion matrix like M but with a modified mass parameter $m' > m$ for a ‘heavier’ fermion. This choice ensures that both J and $J^{-1}K$ have condition numbers lower than K , resulting in a less noisy approximation to the fermion determinant and a corresponding reduction of the simulation cost [3].

An alternative choice of filter is a polynomial $L = P(K)$ of small order p that approximates the inverse K^{-1} , giving fermion action

$$S_{\text{PF}} = \phi_1^\dagger P(K) \phi_1 + \phi_2^\dagger [P(K)K]^{-1} \phi_2. \quad (2.7)$$

This is known as polynomial-filtered HMC [4]. The motivation here is that the polynomial term’s force F_1 is very easy to calculate due to a lack of inverses, and the condition number of the correction term S_2 is reduced as $P(K)K \sim I$. More details on the construction of the corresponding force terms can be found in [Appendix C](#).

This technique can be easily extended to two polynomial filters. If we choose two polynomials $P_1(K)$ and $P_2(K)$ that approximate the inverse such that $Q(K) = P_2(K)/P_1(K)$ is also a polynomial, with order $q = p_2 - p_1$, then we can construct the 2-filter action

$$S_{2\text{PF}} = \phi_1^\dagger P_1(K) \phi_1 + \phi_2^\dagger Q(K) \phi_2 + \phi_3^\dagger [P_2(K)K]^{-1} \phi_3. \quad (2.8)$$

Choosing the polynomials in this way ensures that the force for the intermediate term F_2 is easy to calculate.

The type of polynomial used in this paper is a Chebyshev approximation $P_p(z) \approx 1/z$ of order p , which is parametrized by the real numbers μ and ν . These two parameters are easily chosen such that the net force is minimized whilst the approximation still encompasses K 's eigenvalues; details on this optimization procedure are given in [Appendix A](#). This leaves only the integer parameter p to 'tune'. This compares favourably with mass preconditioning, which has the *real* parameter m' to tune.

2.3. Multi-scale integrators

The primary computational benefit from applying one or more filters via (2.5) to the fermion action arises through the ability to use a multiple time-scale integrator [7], which allows for the evolution of each term on a separate scale.

In order to take advantage of a multiple time-scale integrator, we perform frequency splitting to divide the action into a UV-term and an IR-term $S = S_{UV} + S_{IR}$ [13] where

- S_{UV} captures the high-frequency modes of the system (i.e. large forces) whilst S_{IR} captures the low-frequency modes (small forces).
- F_{UV} is significantly cheaper to calculate than F_{IR} .

The first condition allows one to place the expensive S_{IR} term on a coarser evolution scale without instabilities because of the reduced forces, whilst the second condition means one can place S_{UV} on a finer time-scale with minimal increase in cost. The net effect is to reduce the overall computational cost with minimal loss in acceptance rate.

A good candidate for this technique is polynomial filtering (2.7): the polynomial term $\phi^\dagger P(K)\phi$ captures the high energy modes whilst producing a very cheap force, and can hence act as the UV filter S_{UV} . The preconditioner term $\phi^\dagger J^{-1}\phi$ in mass preconditioning (2.6) works in a similar way; however, there is less direct control over the cost as this depends on the mass m' and typically requires tuning.

This UV/IR prescription can be extended to as many terms as desired. For example, as the gauge action S_G is very cheap, it can be placed on a very fine scale. Hence, for the full 1-filter polynomial-filtered HMC action

$$S = S_G + \phi_1^\dagger P(K)\phi_1 + \phi_2^\dagger [KP(K)]^{-1}\phi_2, \quad (2.9)$$

we choose step-sizes $h_G < h_1 < h_2$.

It is popular [5, 9, 10] to choose step-sizes h_i for each action term S_i such that the average forces F_i are related via

$$F_i h_i = \text{constant}. \quad (2.10)$$

The motive behind this is that a term with a larger force causes correspondingly larger shifts in the Hamiltonian H , so a smaller step-size is required to balance the shifts between the terms and ensure numerical stability. This choice does not necessarily give the optimal parameter set for minimizing the cost, but fine tuning the step-sizes $\{h_i\}$ can be prohibitively expensive in practice.

Large force variances can produce correspondingly large variances in the Hamiltonian H if the step-size is too coarse, which results in low acceptance rates and even exceptional configurations. In particular, experience indicates that the variance is important for the filtered pseudo-fermion correction term, where the size of the force is low but the variance is relatively large. Given the large parameter space in this investigation, we choose the conventional method based on balancing the size of the force terms for simplicity, but here we examine the maximal forces \tilde{F}_i and the step-sizes corresponding to $\tilde{F}_i h_i = \text{constant}$. This often yields a

better acceptance rate (than the absolute value) because it captures some aspects of the variance in the force distributions.

There are more sophisticated methods for step-size tuning. For example, some groups tune the scales by ‘matching’ the tails of the force distributions. Another possibility is to calculate Poisson brackets in order to construct an optimizable approximation to the cost function via the shadow Hamiltonian [16]. However, both these methods are inherently more complex to implement.

Most simulations use a nested leapfrog [7] or a higher-order nested Omelyan integrator [8] as the multi-scale integrator, but this constrains each step-size to evenly divide each coarser step-size. It is possible to construct a generalized multi-scale scheme where no such restrictions exist. The basic idea is to treat the ‘time’ integration steps $T[\epsilon] = (P, U) \rightarrow (P, e^{i\epsilon P} U)$ as advancing a time parameter $\tau \rightarrow \tau + \epsilon$, then superimpose different integration schemes for each action term in terms of τ . This scheme is described in detail in [Appendix B](#).

2.4. Tuning in practice

Each of the filtered actions has a wide range of parameters that can be tuned in order to minimize the computational cost. For example, a good number of 2-flavour Wilson-like simulations use a Hasenbusch filter in their actions [9, 10, 15]

$$\begin{aligned} S &= S_0 + S_1 + S_2 \\ &= S_G + \phi_1^\dagger J^{-1} \phi_1 + \phi_2^\dagger J K^{-1} \phi_2, \end{aligned} \quad (2.11)$$

with each term integrated on a different time-scale. This provides four parameters to tune: m' , $h_0 = h_G$, h_1 and h_2 . However, for physically interesting lattices, generating configurations takes a significantly long time, so the number of trajectories used to tune these parameters should be minimized.

The procedure used in this paper is as follows: first, ‘guess’ some values for the mass preconditioning parameter m' based on m . For each choice, one determines the associated forces $\{F_G, F_1, F_2\}$ from a small number of trajectories, then tunes the step-sizes such that $F_i h_i \approx \text{constant}$. Longer Markov chains are then performed in order to determine the acceptance rate P_{acc} . One then tunes the only free parameter, the coarsest step-size h_2 , such that the desired acceptance rate is reached.

In the case of polynomial filtering, we first tune μ and ν by minimizing the net fermion force as described in [Appendix A](#). Then we treat the polynomial order p like m' in the above procedure. The advantage here is that a good choice of p tends to work well for a wide range of target quark masses m , whereas a good choice of m' depends strongly on m .

2.5. Polynomial-filtered mass-preconditioning

At this point it is pertinent to make some remarks comparing the relative efficacy of polynomial filtering and mass preconditioning.

Mass preconditioning (2.6) works best when the difference between the Hasenbusch mass and the target quark mass $\Delta m = m' - m$ is small, as this implies that $J(m')K^{-1}(m) \simeq I$ and hence the force term is correspondingly reduced. However, when Δm and hence m' is made smaller, the inversion cost to evaluate $J^{-1}\phi$ is increased. At light quark masses, a single Hasenbusch filter is unable to simultaneously satisfy the criteria that the filtered force term F_2 is reduced and the high frequency term F_1 is cheap to evaluate. Due to this, to achieve a computationally efficient frequency-splitting scheme, light quark mass simulations introduce multiple mass preconditioning terms [9–12] that distribute the mass differences across multiple Hasenbusch masses $m < m' < m'' < m''' \dots$. As it is not possible to know *a priori* the inversion cost for a given term, this requires performing simulations to tune the hierarchy of Hasenbusch mass parameters,

which becomes more labour-intensive as more scales are introduced. Previous experience can help guide the choice of parameters, but the extent to which this choice is optimal depends on the ensemble, quark masses and gauge coupling being similar to a past run or another published parameter set.

Meanwhile, the efficacy of polynomial filtering (2.7) depends on two factors: the choice of the polynomial and the spectral range of the matrix whose inverse is being approximated. Specifically, the smaller the spectral range of the matrix K , the smaller the order of the polynomial required to achieve a given accuracy.

In our case, we use a Chebyshev approximation $P(z) \simeq 1/z$ whose roots lie on an ellipse. Choosing the parameters (μ, ν) that determine the ellipse is straightforward: one can simply evaluate the size of the force term while adjusting (μ, ν) and look for a minimum. In practice, one finds that the minimum is relatively shallow and hence fine-tuning of (μ, ν) is not required once a reasonable pair of values has been found.

Once this process has been completed, the only remaining parameter to choose is p , the order of the polynomial approximation. The choice of p allows one to directly determine the cost of the high frequency filter term. As p must be an integer, there is no fine-tuning.

Higher values of p provide a greater reduction in the force for the low frequency correction term F_2 , but correspondingly increase the cost for the filter term F_1 . Hence it is beneficial to make use of multiple polynomial filtering terms (2.8) to introduce additional frequency scales [4]. An advantage of polynomial filtering over mass-preconditioning is that the introduction of an additional scales simply involves choosing another (integer) polynomial order q and hence does not require additional fine-tuning.

Noting that if we had a polynomial of very high order we could approximate the inverse exactly, we can consider the order of the polynomial filter as a means of interpolating between the high and low frequency scales. The effectiveness of polynomial filtering is best in the high frequency regime, associated with high energy scales. As we move to lower frequency scales, the order of polynomial required to capture the dynamics increases significantly and the Chebyshev approximation becomes inefficient when compared with a Krylov-space construction. On the other hand, at low frequency scales mass preconditioning becomes more effective as Δm becomes smaller and hence $J(m')K(m)^{-1} \sim I$.

This observation leads us to propose applying a polynomial filter (or several) to a mass preconditioned fermion action, giving

$$S_{PF-MP} = \phi_1^\dagger P(J) \phi_1 + \phi_2^\dagger [JP(J)]^{-1} \phi_2 + \phi_3^\dagger JK^{-1} \phi_3. \quad (2.12)$$

As $m' > m$, the condition number and hence spectral range of $J(m')$ is reduced in comparison to that of $K(m)$, and hence the accuracy of the polynomial $P(J)$ is better than that of $P(K)$ at a fixed order. The use of short polynomials then provides a good approximation to the high energy fluctuations and is cheap to evaluate, simple to tune and provides direct control over the cost of the highest filtering terms. As the highest energy scales are filtered out using polynomials, the filtered mass preconditioner $P^{-1}(J)J^{-1}$ can be placed on a coarse time scale. Hence, the Hasenbusch mass parameter m' can be chosen such that Δm is small to better reduce the force when evaluating the mass preconditioned quark mass term JK^{-1} . The combined algorithm, which we refer to as polynomial-filtered mass-preconditioned HMC (PF-MP HMC) promises to provide the computational benefit of multiple filters with simpler tuning in comparison to plain mass preconditioning.

3. Results

3.1. Simulation parameters

To study the polynomial-filtered mass-preconditioned algorithm, we first compare polynomial filtering (PF) and mass preconditioning (MP) separately to provide a baseline, then we investigate several variants of the combined PF-MP filtering scheme. We use a modified version of the BQCD program [17] to thermalize a small $16^3 \times 32$ lattice with $n_f = 2$ Wilson fermions at $\kappa = 0.15825$, giving pion mass $m_\pi \sim 400$ MeV. The

gauge coupling is $\beta = 5.6$, providing a lattice spacing of $a \sim 0.08$ fm [9]. This is thermalized with 1000 trajectories of length $\tau = 1$, using two Hasenbusch filters. The choice of integrator for all runs is the second-order minimal norm integrator (B.6) under a generalized multi-scale scheme (see Appendix B). See Table 3.1 for more parameters. Note that we use the conjugate gradient algorithm to invert our fermion matrix: this works well with polynomial filtering, which benefits from the use of a multi-shift conjugate gradient algorithm (see Appendix C). More advanced solvers are available and in use elsewhere [18], from which the PF, MP and PF-MP algorithms could equally benefit.

Parameter	Value
Lattice extent	$16^3 \times 32$
Gauge action	Wilson
Fermion action	Even-odd Wilson
Solver	Conjugate gradient
β	5.6
κ	0.15825

Table 3.1: Table of configuration parameters

A machine-independent indicator of the cost of generating independent configurations is the number of K (and J) multiplications N_{mat} required to generate each configuration. However, we also have to take the acceptance rate P_{acc} into account, because if only a few trajectories are accepted it will take many more tries to generate independent configurations. Thus, we use cost function

$$C = N_{\text{mat}}/P_{\text{acc}} \quad (3.1)$$

as a measure of the expense to produce independent configurations. Our choice of cost function has been used before [14].

Throughout this paper, we attempt to tune the acceptance rate to the range $P_{\text{acc}} = [0.65, 0.75]$, as this has been shown to be cost effective for a second order integrator [19]. The quantities we calculate in the following results have errors given by [20]

$$\sigma = \sqrt{2\tau_{\text{int}}} \frac{\sigma_p}{\sqrt{N}} \quad (3.2)$$

where σ_p is the sample standard deviation of the quantity of interest, N is the sample size, and τ_{int} is the integrated autocorrelation time. For the purposes of (3.2), we use the estimate $\tau_{\text{int}} = 12.5$ for all runs, as tests with the plaquette show that the autocorrelation time for each run lies in the range 8 – 15.

3.2. Comparison of polynomial and mass filtering

We begin our analysis by measuring the performance of polynomial filtering (2.7) relative to mass preconditioning (2.6). This will provide a baseline with which we can compare the combined PF-MP algorithm.

Starting with the simplest case of a single filter, there are still several parameters to tune. The mass preconditioned action (1MP) has the free parameter m' , which, for Wilson fermions as considered here, is equivalent to the hopping parameter $\kappa' < \kappa$. As for polynomial filtering, we select the ellipse parameters $(\mu, \nu) = (1.2, 0.9)$ to minimize the force term (see Appendix A) and hold these values fixed throughout the paper. The polynomial action (1PF) then has only one free parameter: p , the polynomial order. Finally, as we are using a multi-scale integrator, both actions have 3 step-sizes $\{h_0 = h_G, h_1, h_2\}$ to tune.

We tune the parameters as described in section 2.4: a set of appropriate κ' and p are chosen, then the step-sizes $\{h_G, h_1, h_2\}$ are tuned according to the force (shown in Figure 3.2) via the balancing scheme (2.10). However, since the gauge term S_G is very cheap to calculate, it is easier to set the step-size h_G to be sufficiently small such that the produced acceptance rates do not vary, then neglect any further tuning. The generalized multi-scale integration scheme (Appendix B) makes this even easier, as we can keep h_G constant across runs without worrying about whether the other step-sizes are multiples. The resulting parameter choices are given in Tables 3.2 and 3.3; note that we express the step-sizes in terms of the number of steps n_j at each scale, which are related to h_j via $h_j = \tau/n_j$. We also show the average number of K (and J) multiplications required to evaluate the forces as a basis for comparison between the two methods.

p	μ	ν	n_2	n_1	n_0	mat/ F_1	mat/ F_2
4	1.2	0.9	48	120	480	6	672(12)
10			36	160	480	18	741(13)
20			24	240	480	38	693(12)

Table 3.2: Single polynomial filter parameters. ‘mat/ F_i ’ denotes the average number of matrix multiplications by K to evaluate the force F_i . There is no inversion required for F_1 , so the number of matrix multiplications needed is exactly $2p - 2$ (see Appendix C).

κ'	n_2	n_1	n_0	mat/ F_1	mat/ F_2
0.154	8	120	480	84.2(4)	677(12)
0.1545	7	96	480	113.8(8)	627(14)
0.155	7	120	480	112.2(8)	631(10)
0.1555	6	120	480	135.4(1.0)	696(12)
0.156	5	120	480	172.7(1.8)	686(14)

Table 3.3: Single mass filter parameters

Figure 3.1 shows the cost C (3.1) for generating each trajectory for the mass preconditioned and the polynomial filtered actions respectively. Looking at this figure, we see that a single mass filter provides a better overall performance than a single polynomial filter, with a cost of $C = 43,800 \pm 3,500$ at $\kappa' = 0.1545$ compared with $C = 87,500 \pm 7,400$ at $p = 10$.

Given that the cost to evaluate the filter term F_1 is significantly less for the polynomial filter (Table 3.2) than for the mass filter (Table 3.3), it is worthwhile to try to further understand the difference between the two filters. We can do this by considering the force terms. Examining Figure 3.2, we see that the force for the filter term F_1 is similar for both cases. However, the average and maximal forces for the correction term F_2/\tilde{F}_2 are much larger in the polynomial case than in the mass preconditioning case. This leads to more molecular dynamics steps n_2 via (2.10) for PFHMC (see Table 3.2), and is the main reason for the higher cost. As shown in Table 3.2 and indicated by the squares in the right-hand graph of Figure 3.1, increasing the polynomial order to reduce this force simultaneously increases the cost to calculate F_1 , making polynomials of very large order inefficient filters.

The results for a single filter term stand to reason. Given that the Hasenbusch filter is constructing a Krylov-space polynomial to approximate the inverse, a short polynomial term of order 10 cannot capture as much of the dynamics as a Hasenbusch filter that requires 80 or more iterations to invert.

As was done in the original polynomial filtering paper [4], we can factor a higher-order polynomial filter into two terms (see (2.8)) without introducing any additional fine tuning. We denote this technique 2PF for brevity. We set the factoring polynomial’s order to $p_1 = 4$ to keep the cost of F_1 low, then vary the




Figure 3.1: Cost function for 1-filter actions. Squares = matrix operations to construct F_1 , triangles = F_2 construction, filled circles = total. There are some extra matrix multiplications required to initialize the pseudo-fermions ϕ and to construct the fermion action S_F given ϕ and U , but these are negligible for all the actions considered in this paper.




Figure 3.2: 1-filter forces. The left hand plots show the mass preconditioned action's forces whilst the right hand plots show the polynomial filtered action's forces. For each fermion term S_1, S_2 , the maximal and average forces are plotted for each choice of κ'/p .

order of the factored polynomial p_2 . The parameter set is shown in Table 3.4. The cost function for 2PF is shown in Figure 3.3 alongside 1MP for comparison. The minimum of $C = 47,700 \pm 3,700$ here is a marked improvement over 1PF's minimum of $C = 87,500 \pm 7,400$, and is quite comparable to 1MP's performance.

p_1	p_2	μ	ν	n_3	n_2	n_1	n_0
4	24	1.2	0.9	24	20	108	480
	34			20	16	80	480
	54			16	30	120	480

Table 3.4: Configuration parameters for 2PF




Figure 3.3: Cost function for 1MP versus 2PF. Squares = matrix operations to construct F_1 , triangles = F_2 construction, empty circles = F_3 construction, filled circles = total.

3.3. Polynomial-filtered mass-preconditioning

The results of the previous section promote the idea of combining of polynomial filtering with mass preconditioning (2.12), where a cheap polynomial filter is placed on top of a Hasenbusch filter. The PF-MP filtering scheme forms a hierarchy. As $\kappa' < \kappa$, the spectral range of the Hasenbusch filter $J(\kappa')$ is smaller than $K(\kappa)$, increasing the accuracy with which a short polynomial filter can approximate the inverse. As was argued in Section 2.5, applying a polynomial filter to $J(\kappa')$ enables one to reduce the mass difference $\Delta\kappa = \kappa - \kappa'$, increasing the effectiveness of the mass preconditioner. By combining the two schemes in this way we get the best of both worlds: the polynomial term provides a cheap high frequency filter while the Hasenbusch term acts to significantly reduce the force variance in the correction term S_3 .

In modern simulations, the use of an action with two mass preconditioners,

$$S_{2MP} = \phi_1^\dagger J_1^{-1} \phi_1 + \phi_2^\dagger J_1 J_2^{-1} \phi_2 + \phi_3^\dagger J_2 K^{-1} \phi_3, \quad (3.3)$$

is common, and we use this as our benchmark to test the PF-MP scheme. Both actions have two parameters to tune. For the 2MP action we have the Hasenbusch filters $J_1(\kappa_1)$ and $J_2(\kappa_2)$, with $\kappa_1 < \kappa_2 < \kappa$. For the PF-MP action we have the order p of the polynomial term $P(J)$ and the mass $\kappa' < \kappa$ of the Hasenbusch term $J(\kappa')$. The cheapest filter in each case was fixed — $\kappa_1 = 0.145$ for 2MP and $p = 4$ for PF-MP — and optimization took place through the choice of intermediate filter κ_2 (for 2MP) or κ' (for PF-MP) and the choice of step-sizes $\{h_0, h_1, h_2, h_3\}$. As in the previous section, we tune the step-size ratios such that $F_i h_i \approx \text{constant}$, then

κ_1	κ_2	n_3	n_2	n_1	n_0
0.145	0.154	8	15	120	480
	0.155	7	20	96	480
	0.1555	6	20	96	480
	0.156	5	20	120	480
	0.1565	4	20	120	480

Table 3.5: Configuration parameters for 2MP

p	μ	ν	κ'	n_3	n_2	n_1	n_0
4	1.2	0.9	0.154	8	20	80	480
			0.155	6	20	120	480
			0.1555	5	20	120	480
			0.156	5	30	120	480
			0.1565	4	30	120	480

Table 3.6: Configuration parameters for PF-MP

tune the coarsest step-size h_3 to the correct acceptance rate. The full range of parameters are detailed in Tables 3.5 and 3.6.

Figure 3.4 shows the forces for the 2MP and PF-MP runs. Whereas for the single-filter actions (Section 3.2) the correction term for polynomial filtering has a much greater force variance than that for mass preconditioning, here, the corresponding polynomial correction term S_2 for PF-MP has a maximal force only slightly larger than that of the Hasenbusch correction term S_2 for 2MP. This supports the prior argument that polynomial filtering (at a fixed order) is more effective on $J(\kappa')$ than on $K(\kappa)$; we are filtering at a heavier mass $\kappa' < \kappa$, with an associated suppression in the long range physics.

Figures 3.5, 3.6 and 3.7 show the matrix operation count, acceptance rate, and cost respectively, with 2MP on the left and PF-MP on the right. Looking at Figure 3.7, the optimal point for 2MP is at $\kappa_2 = 0.1555$ with cost $C = 31,000 \pm 2,200$, whereas for PF-MP it is at $\kappa' = 0.155$ with cost $C = 29,000 \pm 1,800$. We see that the PF-MP scheme can perform just as well as mass preconditioning in this instance.

3.4. Tests with 3-level filters

We have examined the PF-MP action in the case of a single polynomial filter applied to a single mass preconditioner, which we can denote as 1PF-1MP. Within the PF-MP scheme, as for plain polynomial filtering, we can increase the order of the polynomial filter and then factor that into two terms to see if the introduction of an additional intermediate scale provides any additional benefit. This does not require any additional fine tuning, as the choice of polynomial order p provides direct control over the cost and scale of the filter terms, independent of the quark mass. We denote the scheme with a 2-level polynomial filter and a single mass preconditioner as 2PF-1MP:

$$\begin{aligned}
S_{2PF-1MP} = & \phi_1^\dagger P_1(J) \phi_1 + \phi_2^\dagger Q(J) \phi_2 \\
& + \phi_3^\dagger [JP_2(J)]^{-1} \phi_3 + \phi_4^\dagger JK^{-1} \phi_4.
\end{aligned} \tag{3.4}$$




Figure 3.4: 2-filter forces. The left hand plots show the 2MP forces whilst the right hand plots show the PF-MP forces. For each fermion term S_1, S_2, S_3 , the maximal and average forces are plotted for each choice of κ_2/κ' . The third force F_3 is the same in each case because the third action term S_3 is also the same.




Figure 3.5: Matrix operation counts for 2-filter actions. The squares = matrix ops from constructing F_1 , triangles = F_2 construction, empty circles = F_3 construction, filled circles = total. The errors are omitted here due to being smaller than the marker size.




Figure 3.6: Acceptance rates for 2-filter actions




Figure 3.7: Cost function for 2-filter actions. The squares = matrix ops from constructing F_1 , triangle = F_2 construction, empty circles = F_3 construction, filled circles = total.

For completeness we also examine the 1PF-2MP scheme with a single polynomial filter and 2 levels of mass preconditioning,

$$S_{1PF-2MP} = \phi_1^\dagger P(J_1) \phi_1 + \phi_2^\dagger [J_1 P(J_1)]^{-1} \phi_2 + \phi_3^\dagger J_1 J_2^{-1} \phi_3 + \phi_4^\dagger J_2 K^{-1} \phi_4; \quad (3.5)$$

however, this does introduce an additional mass parameter that requires fine tuning.

For the 1PF-2MP scheme, we fix the polynomial order at $p = 4$ as with PF-MP, and set κ_1 to 0.145 to match the 2MP runs. For the 2PF-1MP scheme, we choose $p = p_2 = 24$, factored into terms of order $p_1 = 4$ and $q = p_2 - p_1 = 20$, leaving only the single Hasenbusch parameter κ' to tune. See Tables 3.7 and 3.8 for a full list of parameters. The forces for 1PF-2MP and 2PF-1MP are shown in Figure 3.8; note that the forces associated with the S_3 term are significantly smaller for 2PF-1MP than for 1PF-2MP. Figures 3.9, 3.10 and 3.11 show the matrix operation count, acceptance rate and cost respectively.

p	μ	ν	κ_1	κ_2	n_4	n_3	n_2	n_1	n_0
4	1.2	0.9	0.145	0.153	11	12	16	96	480
				0.154	8	15	15	96	480
				0.1555	6	20	20	96	480
				0.1565	4	24	20	96	480

Table 3.7: Configuration parameters for 1PF-2MP

p_1	p_2	μ	ν	κ'	n_4	n_3	n_2	n_1	n_0
4	24	1.2	0.9	0.153	10	5	16	80	480
				0.154	9	6	24	120	480
				0.1555	6	8	20	120	480
				0.1565	4	10	24	120	480

Table 3.8: Configuration parameters for 2PF-1MP

For ease of comparison, the cost function for all the actions considered in this paper are presented in Figure 3.12, aside from 1PF which has a significantly higher cost than the other actions. Looking at this figure, the three PF-MP schemes all have a similar cost minimum, which is as good as or better than the 2MP benchmark. More important is the relative dependence on the free mass parameter, κ' or κ_2 . We can see that for 2MP, 1PF-1MP and 1PF-2MP that a poor choice of κ'/κ_2 can lead to a significant increase in the cost function (see e.g. $\kappa'/\kappa_2 = 0.1565$), whereas the 2PF-1MP cost function has only a very weak dependence on the Hasenbusch mass parameter. This demonstrates that no fine tuning of κ' is required for 2PF-1MP to achieve optimal performance.




Figure 3.8: 3-filter forces. The left hand plots show the 1PF-2MP forces whilst the right hand plots show the 2PF-1MP forces. For each fermion term, the maximal and average forces are plotted for each choice of κ_2/κ' .




Figure 3.9: Matrix operation counts for 3-filter actions. The squares = matrix ops due to constructing the force F_1 , triangles = F_2 construction, empty circles = F_3 construction, diamonds = F_4 construction, filled circles = total.




Figure 3.10: Acceptance rates for 3-filter actions




Figure 3.11: Cost function for 3-filter actions. The squares = matrix ops due to constructing the force F_1 , triangles = F_2 construction, empty circles = F_3 construction, diamonds = F_4 construction, filled circles = total.




Figure 3.12: Cost function (3.1) for this paper's actions.

4. Conclusion

We have compared the polynomial filtered and mass preconditioned HMC algorithms, and found that a 2-level polynomial filter provides a benefit similar to a single mass preconditioner. We proposed combining the two methods to provide a multi-level frequency-splitting scheme with minimal fine tuning of the action parameters. This was partly motivated by noting that the values (μ, ν) determining the Chebyshev polynomial roots produce a shallow minimum in the polynomial force term, and hence do not need fine tuning, leaving the polynomial order p as the only free parameter.

Any form of Sexton–Weingarten integration with a large number of terms requires a sensible choice of the relative time scales to achieve good performance. The tuning of the different time steps for our study of multi-level algorithms was aided by using a generalized multi-scale integration scheme, permitting any choice of step-size for each action term. This made it simple to use the force balancing method ' $F_i h_i = \text{constant}$ ' to select the scale for each action term based on its (maximal or average) force.

The polynomial-filtered mass-preconditioned (PF-MP) algorithm was investigated with $n_f = 2$ flavours of dynamical quarks, using several different combinations of polynomial and Hasenbusch filters, and compared to 2-level mass preconditioning (2MP) as a baseline. We found that the 2PF-1MP action yielded a cost function that was as good as or better than the 2MP action, with a significant reduction in the tuning effort required to optimize the overall cost. The 2MP action has two real Hasenbusch parameters κ_1, κ_2 that need to be tuned. In contrast, the 2PF-1MP action did not need any fine tuning: it showed almost no dependence on the Hasenbusch parameter κ' , and the orders of the polynomial terms (as integers) were easily chosen to optimize the cost.

This study was performed at an intermediate quark mass $m_\pi \sim 400$ MeV as a proof of the viability of the PF-MP scheme. Simulations at lighter quark masses typically introduce additional filters to further

ameliorate the cost of these simulations, with some groups even using 6-level mass preconditioning [12]. At these light quark masses, the PF-MP algorithm can potentially provide an easier path to gain the benefits of multi-level frequency splitting.

Acknowledgements

The authors would like to thank M. Peardon for valuable discussions. We have used a modified version of BQCD [17] to generate the configurations in this work. This work was supported by the Victorian Life Sciences Computation Initiative (VLSCI), an initiative of the Victorian Government, Australia, on its Facility hosted at the University of Melbourne, Grant Number NCE31. This work was also supported by resources provided by The Pawsey Supercomputing Centre with funding from the Australian Government and the Government of Western Australia. This investigation was supported by the Australian Research Council under Grant Numbers FT100100005, DP120104627, DP140103067, and DP150103164.

Appendix A. Chebyshev polynomials

The Chebyshev approximation to the inverse K^{-1} takes the form

$$K^{-1} \approx P_n(K) = a_n \prod_{k=1}^n (K - z_k), \quad (\text{A.1})$$

where the roots z_k are defined via

$$z_k = \mu(1 - \cos \theta_k) - i\sqrt{\mu^2 - \nu^2} \sin \theta_k, \quad \theta_k = \frac{2\pi k}{n+1} \quad (\text{A.2})$$

and the normalization a_n is given by

$$a_n = \frac{1}{\mu \prod_{k=1}^n (\mu - z_k)}. \quad (\text{A.3})$$

This has three free parameters — n , μ , ν — that can be adjusted to suit the fermion matrix K we wish to approximate the inverse of.

The roots describe an ellipse in complex space which passes through the origin, with semi-major axis along the positive real line with length μ and semi-minor axis $\sqrt{\mu^2 - \nu^2}$. If we add the origin to the roots to make a set of $n+1$ points, these points are distributed at equal angles around the ellipse. See Figure A.1 for an example. The approximation is only effective at points within this ellipse, so one should choose $\mu \geq \nu > 0$ such that the spectrum of K is contained. For the lattice configuration considered in this paper, the eigenvalues of K go from $\lambda_{\min} = 3.2 \times 10^{-5}$ to $\lambda_{\max} = 2.2$. This means we must choose $\mu > 2.2/2 = 1.1$ for a good approximation.

Aside from needing to bound the eigenvalues, we have a lot of freedom in the choice of (μ, ν) . The tuning procedure taken in this work is to simply choose the set (μ, ν) that minimizes the average force. On the configuration used in this paper, we found that $\mu = 1.2$ and $\nu = 0.9$ gave the best forces across our different choices for n and (for the PF-MP actions) mass parameter κ' .

The remaining parameter, the polynomial order n , can then be varied to ensure a good hierarchy of forces in the fermion action. This is similar to choosing κ' in mass preconditioning; however, as n must be an integer, the need for fine-tuning is excluded. Another advantage of polynomial filtering is that choosing a particular n will filter out a similar proportion of the action no matter the choice of mass parameter κ , whereas, for mass preconditioning, κ' has to be varied to find a particular splitting.




Figure A.1: The roots of a 4th order Chebyshev polynomial approximation to K^{-1} . The roots are shown as circles, and if the origin (cross) is included they are evenly distributed around an ellipse.

A useful property of the Chebyshev polynomial filters is that two approximations P_{p_1}, P_{p_2} with $p_2 > p_1$ and the same (μ, ν) factorize if $(p_1 + 1)$ divides $(p_2 + 1)$. The ratio Q is then a polynomial of order $q = p_2 - p_1$, and can be used as an intermediate filter via (2.8), as was shown in previous work [4].

Appendix B. Generalized multi-scale integrator

The generalized multi-scale integrator presented here is an extension of the generalized leapfrog integrator described in [4] and is mentioned in [21]. The general idea is as follows: assume that for each action term S_i we have an integration scheme that preserves a Hamiltonian $H_i = T + S_i$ through a series of ‘time’ \hat{T} (B.4b) and ‘space’ \hat{S}_i (B.4a) updates. Supposing all our step-sizes are positive, we can treat the series of time updates like they advance a time parameter τ from 0 to h . The generalized multi-scale integrator for the full Hamiltonian $H = T + \sum_i S_i$ then works by advancing through this time, inserting the action term updates (\hat{S}_i) at times corresponding to their position in the original integrators. To clarify this process, an example is depicted in Figure B.2, where a 3-step leapfrog integrator and a 1-step second-order minimal norm integrator are combined.

The purpose of the rest of this appendix is to give a concrete definition of the generalized multi-scale integrator scheme and to prove that, with the right component integrators, it is time-reversible and area-preserving as required.

Appendix B.1. Basics

An integration step for HMC takes the system from some state (P, U) to another state (P', U') . In this section, we will use the notation \hat{M} for an integration step, with

$$(P', U') = \hat{M}(P, U). \quad (\text{B.1})$$

Integration steps are typically parametrized by some step-size ϵ and we denote this with $\hat{M}[\epsilon]$. It is useful to note that the set of all deterministic integration steps forms a group under composition.

In order for HMC to produce configurations that follow the desired probability distribution $\exp[-S]$, the integration scheme \hat{M} used must be

- time-reversible:

$$\hat{M}[-\epsilon](-P', U') = (P, U),$$

but as the momentum P only enters the kinetic term T quadratically, we can ignore the minus sign on P' and write

$$\begin{aligned} \hat{M}[-\epsilon](P', U') &= (P, U) \\ \Rightarrow \hat{M}[-\epsilon] &= \hat{M}^{-1}[\epsilon] \end{aligned} \quad (\text{B.2})$$

- area preserving:

$$\det \frac{\partial(P', U')}{\partial(P, U)} = 1 \quad (\text{B.3})$$

Our atomic steps for constructing an appropriate integration scheme come from Hamilton's equations, and have two flavours:

$$\hat{S}[\epsilon] : \hat{S}[\epsilon](P, U) = (P - \epsilon F(U), U) \quad (\text{B.4a})$$

and

$$\hat{T}[\epsilon] : \hat{T}[\epsilon](P, U) = (P, e^{i\epsilon P} U), \quad (\text{B.4b})$$

where $F(U) = \left. \frac{\partial S}{\partial U} \right|_U$ is the force term. When we have multiple action terms $S = S_1 + S_2 + \dots$, we can use integration steps \hat{S}_i for each force term F_i . These atomic steps are both time-reversible and area preserving.

We denote a scheme composed solely of \hat{T} and \hat{S}_i steps *symplectic*, as each step is tangential to the curve in phase space where the Hamiltonian H is preserved.

Appendix B.2. Area preservation

When an integration scheme is composed of several steps, it is easy to prove area-preservation: since $\det AB = \det A \det B$, any product of area-preserving steps (such as \hat{S}_i and \hat{T}) is automatically area-preserving. In particular, symplectic schemes are area-preserving.

Appendix B.3. Time reversibility

\hat{S} and \hat{T} have the special property that

$$\begin{aligned} \hat{S}[a+b] &= \hat{S}[a]\hat{S}[b], \\ \hat{T}[a+b] &= \hat{T}[a]\hat{T}[b], \end{aligned}$$

for all $a, b \in \mathbb{R}$ and $\hat{S}[0] = \hat{T}[0] = \hat{I}$. Note that this implies the property $[\hat{S}[a], \hat{S}[b]] = 0 = [\hat{T}[a], \hat{T}[b]]$, so the sets $\{\hat{S}[\epsilon]\}$ and $\{\hat{T}[\epsilon]\}$ form Abelian groups.

Due to the above property, a symplectic integration scheme can be written in the form

$$\hat{M}[h] = \hat{T}[a_{n+1}] \mathcal{L} \prod_{i=1}^n (\hat{S}[b_i] \hat{T}[a_i]), \quad (\text{B.5})$$

where $a_i, b_i \in \mathbb{R}$, $\mathcal{L} \prod$ denotes a product with the first index acting first (i.e. on the right) and $\sum a_i = \sum b_i = h$ by convention. To ensure this expression is unique, we mandate that $b_i \neq 0 \ \forall i = 1, \dots, n$ and $a_i \neq 0 \ \forall i = 2, \dots, n$.

Theorem 1. *Given a set of time-reversible steps \hat{A}_i , the integration scheme*

$$\hat{M} = \mathcal{L} \prod_{i=1}^n \hat{A}_i$$

is time-reversible if $\hat{A}_i = \hat{A}_{n-i+1} \ \forall i$. We denote such an integration scheme symmetric.

Theorem 2 (Corollary). *A symplectic integration scheme*

$$\hat{M}[h] = \hat{T}[a_{n+1}] \mathcal{L} \prod_{i=1}^n (\hat{S}[b_i] \hat{T}[a_i]),$$

is time-reversible if

$$\begin{aligned} a_i &= a_{n-i+2}, \\ b_i &= b_{n-i+1}. \end{aligned}$$

This allows one to determine whether a given integration scheme is time-reversible by trying to write it as a symmetric product of time-reversible steps. For example, the second order minimal-norm space-time-space integration step is

$$\hat{M}_{2\text{MNSTS}}(\epsilon) = \hat{S}[\lambda\epsilon] \hat{T}\left[\frac{\epsilon}{2}\right] \hat{S}[(1-2\lambda)\epsilon] \hat{T}\left[\frac{\epsilon}{2}\right] \hat{S}[\lambda\epsilon], \quad (\text{B.6})$$

and this is time-reversible by virtue of being symmetric (Theorem 1).

In order to easily generalize to multi-step schemes, we note that if \hat{M} is time-reversible, then so is \hat{M}^n . This means that, for example, an n -step second-order minimal-norm space-time-space integration scheme is time-reversible and area preserving (due to being symplectic).

Appendix B.4. The generalized multi-scale integrator

To define the generalized multi-scale integrator, we need to introduce a new operator.

Consider integrating a Hamiltonian $H = T + \sum_i S_i$ with several action terms. Note that regardless of how many action terms S_i we have, we only ever have one kind of ‘time’ update \hat{T} . Thus, supposing we only integrate in one direction with all $\epsilon \geq 0$ or $\epsilon \leq 0$, it makes sense to parametrize the progress of $\hat{T}[\epsilon]$ updates via a time parameter μ that ranges from 0 to h . This time parameter can be attached to the force updates

$$\hat{M}[\epsilon] \rightarrow \hat{M}[\epsilon, \mu = \tau], \quad (\text{B.7})$$

which does not affect the action of the integration step, but it does allow one to define a useful operator:

Definition 1. The *time-step insertion operator* \mathcal{T}_A^B acting on some product of operators with assigned time parameters is defined as

$$\begin{aligned} \mathcal{T}_A^B \prod_{i=1}^n \hat{M}_i[\epsilon_i, \mu = \tau_i] \\ \equiv \hat{T}[\tau'_{n+1} - \tau'_n] \mathcal{L} \prod_{i=1}^n (\hat{M}_i[\epsilon_i] \hat{T}[\tau'_i - \tau'_{i-1}]) \end{aligned} \quad (\text{B.8})$$

where the time parameters have been reordered $\{\tau_i\} = \{\tau'_i\}$ such that $\tau'_i \leq \tau'_{i+1} \forall i = 1, \dots, n$, and we define $\tau'_0 = A$ and $\tau'_{n+1} = B$. Typically, we choose $A = 0$ and $B = h$, and this will be written as $\mathcal{T} = \mathcal{T}_0^h$. See Figure B.1 for a depiction of this operator in action.

A symplectic scheme (B.5) can thus be written as

$$\hat{M}[h] = \mathcal{T} \prod_{i=1}^n \hat{S}[b_i, \mu = c_i] \quad (\text{B.9})$$




Figure B.1: The time-step insertion operator \mathcal{T} . It works by inserting time steps \hat{T} between each space step update \hat{S} and the given temporal endpoints. Here, we depict some space steps as vectors in the upward direction based at particular points on temporal axis τ (top), then apply the time insertion operator by inserting time steps along the horizontal (bottom).

where $c_i = \sum_{j=1}^i a_j$.

This form of the symplectic integration scheme allows one to easily define the generalized multi-scale integrator:

Definition 2. Consider a Hamiltonian $H = T + \sum_i S_i$ with several action terms. Suppose that for each term S_i we have a symplectic integration scheme

$$\hat{M}_i[h] = \mathcal{T} \prod_{j=1}^{n_i} \hat{S}_i[b_j^{(i)}, \mu = c_j^{(i)}]$$

that preserves $H_i = T + S_i$. Then the *generalized multi-scale integrator* for the full Hamiltonian H is given by

$$\hat{M}_{\text{gen}}[h] = \mathcal{T} \left(\prod_{i=1} \prod_{j=1}^{n_i} \hat{S}_i[b_j^{(i)}, \mu = c_j^{(i)}] \right). \quad (\text{B.10})$$

This construction is unambiguous: if there is a $c_i^{(k)} = c_j^{(l)}$, the order is insignificant since $[\hat{S}_k, \hat{S}_l] = 0$. Also note that \hat{M}_{gen} is symplectic by construction, and hence area-preserving.

Appendix B.4.1. Reversibility

In order to show that the generalized multi-scale integrator is time-reversible, it is necessary to determine what reversibility looks like for a time-step inserted product of operators.

Theorem 3. Suppose we have an integration scheme

$$\hat{M} = \mathcal{T}_A^B \prod_{i=1}^n \hat{A}_i[\mu = b_i]. \quad (\text{B.11})$$

where $[\hat{A}_i, \hat{A}_j] = 0$ if $b_i = b_j$. Then \hat{M} is symmetric (and hence reversible) iff for every operator $\hat{A}[\mu = b]$ in the product, there is also an operator of the form $\hat{A}[\mu = A + B - b]$.

Proof. Under the action of the time-step insertion operator, we can rearrange the operators \hat{A}_i such that $b_i \leq b_{i+1}$. Expanding the time-step insertion operator then gives:

$$\hat{M} = \hat{T}[b_{n+1} - b_n] \mathcal{L} \prod_{i=1}^n (\hat{A}_i \hat{T}[b_i - b_{i-1}]),$$

where $b_0 = A$ and $b_{n+1} = B$. Using Theorem 2, this is symmetric iff

$$\hat{A}_i = \hat{A}_{n-i+1} \quad (\text{B.12})$$

and

$$b_i - b_{i-1} = b_{n-i+2} - b_{n-i+1} \quad \forall i = 0, \dots, n+1$$

Rearranging the second condition gives

$$\begin{aligned} b_i + b_{n-i+1} &= b_{i-1} + b_{n-i+2} \\ &= b_{i-2} + b_{n-i+3} \\ &= \dots \\ &= b_0 + b_{n+1} = A + B \end{aligned}$$

so

$$b_i + b_{n-i+1} = A + B \quad \forall i = 0, \dots, n+1 \quad (\text{B.13})$$

The two conditions (B.12) and (B.13) together are equivalent to saying that for each the operator $\hat{A}_i[\mu = b_i]$ in the product, we also have $\hat{A}_{n-i+1}[\mu = b_{n-i+1}] = \hat{A}_i[\mu = A + B - b_i]$. \square

Theorem 4. *If the constituent symplectic integrators of the generalized multi-scale integrator are symmetric, then the generalized multi-scale integrator is also symmetric and hence time-reversible.*

Proof. Consider the generalized multi-scale integrator (B.10) with two schemes:

$$\hat{M}_{\text{gen}}[h] = \mathcal{T} \left(\prod_{i=1}^{n_1} \hat{S}_1[b_i^1, \mu = c_i^1] \prod_{i=1}^{n_2} \hat{S}_2[b_i^2, \mu = c_i^2] \right)$$

By assumption, the two constituent schemes are symmetric. Hence, by Theorem 3, it follows that each $\hat{S}_1[b, \mu = c]$ has a mirror $\hat{S}_1[b, \mu = h - c]$ and each $\hat{S}_2[b, \mu = c]$ has a mirror $\hat{S}_2[b, \mu = h - c]$.

But that means *every* operator $\hat{M}[b, \mu = c]$ in the product has a mirror $\hat{M}[b, \mu = h - c]$. Hence, by Theorem 3, \hat{M}_{gen} is symmetric and reversible. This extends trivially to an arbitrary number of schemes. \square

Appendix B.5. Implementation

Appendix B.5.1. Example algorithm

The integration scheme for an individual action term S_i can be expressed as two arrays: `T_steps` which holds the time \hat{T} updates and `S_steps` which holds the space S_i updates, ordered such that the scheme can be enacted by a simple loop:

```

for i in (1,length(T_steps)):
    integrate_T(step=T_steps[i])
    integrate_S_i(step=S_steps[i])

```

For example, a 2-step leapfrog algorithm

$$\hat{A}[h] = \hat{S}[h/4]\hat{T}[h/2]\hat{S}[h/4]\hat{T}[h/2]\hat{S}[h/4] \quad (\text{B.14})$$

can be expressed as

```

T_steps = (0, h/2, h/2)
S_steps = (h/4, h/2, h/4)

```

To implement the generalized multi-scale integrator for an action with n terms, we can combine the $2n$ arrays as follows:

```

n = <number of S_i action terms>
T_steps_i = [<T steps for int method 1>, ...]
S_steps_i = [<S steps for int method 1>, ...]

pop() = remove first element of the array

new_T_steps = []
new_Si_steps = [[]]
tau = 0
d_tau = 0
while tau < traj_length {
    # Find the smallest time step
    # out of the potential next ones
    d_tau = min(T_steps_i[:,1])
    # Add the new time step
    new_T_steps.append(d_tau)
    for i in (1,n) {
        # If it is time to insert a 'S_i' step
        if d_tau == T_steps_i[i][1] {
            # Add the next space step to the new list
            new_Si_steps[i].append(S_steps_i[i][1])
            # Remove the time and space steps
            # from the old lists
            S_steps_i.pop()
            T_steps_i[i].pop()
        } else {
            # Add a 'do nothing' step to the new list
            new_Si_steps[i].append(0)
            # Decrement the next time step
            # (moving forward in time)
            T_steps_i[i][1] -= d_tau
        }
    }
}

```

```

}
# ! At this stage, each individual scheme will
# ! be up to time t = tau + d_tau
tau += d_tau
}

```

The generalized multi-scale integrator can then be enacted via

```

for i in (1,length(new_T_steps)):
    integrate_T(step=new_T_steps[i])
    for j in (1,n):
        integrate_S_j(step=new_Si_steps[j][i])

```

Note that the order of the action updates \hat{S}_j in the inner loop does not matter since $[\hat{S}_i, \hat{S}_j] = 0$.

Appendix B.5.2. Algorithm demonstration

Suppose we choose to use a 3-step leapfrog and a 1-step second-order minimal norm scheme:

$$\begin{aligned}\hat{M}_1[h] &= \hat{S}_1[h/6]\hat{T}[h/3]\hat{S}_1[h/3]\hat{T}[h/3]\hat{S}_1[h/3]\hat{T}[h/3]\hat{S}_1[h/6], \\ \hat{M}_2[h] &= \hat{S}_2[\lambda h]\hat{T}[h/2]\hat{S}_2[(1-2\lambda)h]\hat{T}[h/2]\hat{S}_2[\lambda h].\end{aligned}$$

These can be written in array form as

$$\begin{aligned}T_1 &= (0, h/3, h/3, h/3), \\ S_1 &= (h/6, h/3, h/3, h/6), \\ T_2 &= (0, h/2, h/2), \\ S_2 &= (\lambda h, (1-2\lambda)h, \lambda h).\end{aligned}$$

Merging these two schemes by hand (see Figure B.2) shows that the resultant scheme should take the form

$$\begin{aligned}T' &= (0, h/3, h/6, h/6, h/3), \\ S'_1 &= (h/6, h/3, 0, h/3, h/6), \\ S'_2 &= (\lambda h, 0, (1-2\lambda)h, 0, \lambda h).\end{aligned}$$

We step through the algorithm in Figure B.3 to show that it indeed produces this result.




Figure B.2: A generalized multi-scale integrator demonstration. Initially, S_1 is integrated with a 3-step leapfrog integrator (top) whilst S_2 uses a 1-step second-order minimal norm (bottom). To combine these two into a generalized multi-scale scheme, the two schemes are overlapped based on the 'time' axis \hat{T} (centre), then the time steps are recalculated based on where each 'space' update \hat{S}_i takes place. Each scheme is also expressed by an array of time steps $\mathbf{T_steps}$ and one or more arrays of space steps $\mathbf{S_steps_i}$, which can be used in code as described in [Appendix B.5.1](#).




Figure B.3: Demonstration of the generalized multi-scale integrator algorithm (Appendix B.5.1). The state of each variable at the start (top) and after each loop iteration is shown on the left, and on the right we have the state of the constructed scheme at these points. With each loop of the algorithm, we find the next point in time τ where we need to insert a space step, move to time τ with \hat{T} , then insert the space steps \hat{S}_i which are at time τ . The algorithm is complete when we reach $\tau = h$ (bottom), giving the correct generalized multi-scale scheme (c.f. Figure B.2).

Appendix B.6. Error terms

A vital consideration for a generalized multi-scale integrator is how its error terms compare to those of the composite integration schemes. For example, the leapfrog space-time-space scheme with step-size h and $n = t/h$ steps has error term

$$\tilde{H}_{LPF} - \hat{H} = h^2 \left(\frac{1}{12} [\hat{S}, [\hat{S}, \hat{T}]] + \frac{1}{24} [\hat{T}, [\hat{S}, \hat{T}]] \right) + O(h^4) \quad (\text{B.15})$$

where \hat{H} is the true Hamiltonian and \tilde{H}_{LPF} is the actual effect of the integrator, and the second order minimal norm space-time-space scheme (B.6) has error term

$$\begin{aligned} \tilde{H}_{2MN} - \hat{H} = h^2 & \left(\frac{6\lambda^2 - 6\lambda + 1}{12} [\hat{S}, [\hat{S}, \hat{T}]] \right. \\ & \left. + \frac{1 - 6\lambda}{24} [\hat{T}, [\hat{S}, \hat{T}]] \right) + O(h^4). \end{aligned} \quad (\text{B.16})$$

Such error terms are usually calculated by recursively applying the Baker–Campbell–Hausdorff formula for a symmetric product

$$\ln(e^{hA} e^{hB} e^{hA}) = h(2A + B) - \frac{h^3}{6} ([B, [A, B]] + [A, [A, B]]) + O(h^5), \quad (\text{B.17})$$

from the centre of a symmetric scheme.

Let us consider the general case of a step in this expansion for a generalized multi-scale integrator with Hamiltonian $\hat{H} = \hat{T} + \sum_i \hat{S}_i$, writing $\hat{T}[a] = e^{ah\hat{T}}$ for this section only. This comes in two flavours:

$$e^{\alpha h \hat{T}} \exp \left[\beta h \hat{T} + \sum_i \gamma_i h \hat{S}_i \right] e^{\alpha h \hat{T}} \quad (\text{B.18a})$$

and

$$e^{\delta_i h \hat{S}_i} \exp \left[\beta h \hat{T} + \sum_i \gamma_i h \hat{S}_i \right] e^{\delta_i h \hat{S}_i}. \quad (\text{B.18b})$$

Using (B.17), these expand to

$$\begin{aligned} \exp & \left[(2\alpha + \beta) h \hat{T} + \sum_i \gamma_i h \hat{S}_i + \frac{\alpha(\alpha + \beta) h^3}{6} \sum_i \gamma_i [\hat{T}, [\hat{S}_i, \hat{T}]] \right. \\ & + \frac{\alpha h^3}{6} \sum_i \gamma_i^2 [\hat{S}_i, [\hat{S}_i, \hat{T}]] \\ & \left. + \frac{\alpha h^3}{3} \sum_i \sum_{j>i} \gamma_i \gamma_j [\hat{S}_i, [\hat{S}_j, \hat{T}]] + O(h^5) \right] \end{aligned}$$

and

$$\begin{aligned} \exp & \left[\beta h \hat{T} + \sum_{j \neq i} \gamma_j h \hat{S}_j + (\gamma_i + 2\delta_i) \hat{S}_i - \frac{\beta^2 \delta_i h^3}{6} [\hat{T}, [\hat{S}_i, \hat{T}]] \right. \\ & - \frac{\beta \delta_i (\delta_i + \gamma_i) h^3}{6} [\hat{S}_i, [\hat{S}_i, \hat{T}]] \\ & \left. - \frac{\beta \delta_i h^3}{6} \sum_{j \neq i} \gamma_j [\hat{S}_i, [\hat{S}_j, \hat{T}]] + O(h^5) \right]. \end{aligned}$$

Note that for any given i , the coefficients for $[\hat{T}, [\hat{S}_i, \hat{T}]]$ and $[\hat{S}_i, [\hat{S}_i, \hat{T}]]$ only involve the coefficients for \hat{T} and \hat{S}_i from the initial expressions (B.18a) and (B.18b). Hence, the resulting coefficients for these terms when expanding a full scheme *must* be the same as what would result with only \hat{T} and \hat{S}_i steps. In the case of a generalized multi-scale scheme, the scheme's construction (B.10) thus ensures that the coefficients of these terms are identical to the ones for the composite integrators, for example (B.16).

The only new terms are the cross terms $[\hat{S}_i, [\hat{S}_j, \hat{T}]]$, $i \neq j$. The coefficients for these terms depend on how both \hat{S}_i and \hat{S}_j are integrated, and hence the cross terms typically reduce the benefit gained by placing one action term or the other on a finer time-scale.

As an example, the integration scheme described in Appendix B.5.2 has error term

$$\begin{aligned} \tilde{H} - \hat{H} = h^2 & \left(\frac{1}{108} [\hat{S}_1, [\hat{S}_1, \hat{T}]] + \frac{1-6\lambda}{24} [\hat{S}_2, [\hat{S}_2, \hat{T}]] \right. \\ & + \frac{1}{216} [\hat{T}, [\hat{S}_1, \hat{T}]] + \frac{6\lambda^2 - 6\lambda + 1}{12} [\hat{T}, [\hat{S}_2, \hat{T}]] \\ & \left. + \frac{1-8\lambda}{36} [\hat{S}_1, [\hat{S}_2, \hat{T}]] \right) + \mathcal{O}(h^4). \end{aligned}$$

Appendix C. Force terms

This appendix describes the force terms for a variety of fermion actions in order to show how they could be implemented in code.

Appendix C.1. Basic HMC

The force term for the basic fermion action

$$S_F = \phi^\dagger K^{-1} \phi \tag{C.1}$$

is

$$F = \frac{\partial S_F}{\partial U} = -\phi^\dagger K^{-1} \frac{\partial K}{\partial U} K^{-1} \phi. \tag{C.2}$$

The form of $\frac{\partial K}{\partial U}$ is dependent on the choice of fermion matrix K , e.g. Wilson or Clover.

Appendix C.2. Polynomial-filtered HMC

Consider the 1-filter PFHMC action

$$S_{\text{1pf}} = \phi_1^\dagger P(K) \phi + \phi_2^\dagger [P(K)K]^{-1} \phi_2. \tag{C.3}$$

Given a polynomial $P(K)$ in the form

$$P(K) = c_n \prod_{i=1}^n (K - z_i),$$

we can write the polynomial force term as

$$\begin{aligned}
F_1 &= \frac{\partial S_1}{\partial U} \\
&= \phi_1^\dagger \frac{\partial}{\partial U} \left[c_n \prod_{i=1}^n (K - z_i) \right] \phi_1 \\
&= \phi_1^\dagger \sum_{i=1}^n \left[c_n \prod_{j=1}^{i-1} (K - z_j) \frac{\partial K}{\partial U} \prod_{j=i+1}^n (K - z_j) \right] \phi_1 \\
&= \sum_{i=1}^n \eta_i^\dagger \frac{\partial K}{\partial U} \chi_i,
\end{aligned} \tag{C.4}$$

where

$$\chi_i = c_n \prod_{j=i+1}^n (K - z_j) \phi_1, \tag{C.5a}$$

$$\eta_i = \prod_{j=1}^{i-1} (K - z_j^*) \phi_1. \tag{C.5b}$$

If we construct these intermediate fields incrementally, the calculation of F_1 only requires $(2n - 2)$ matrix multiplications in addition to the ones required to calculate $\frac{\partial K}{\partial U}$.

As for the correction term F_2 , we have

$$\begin{aligned}
F_2 &= \frac{\partial S_2}{\partial U} \\
&= \phi_2^\dagger \frac{\partial}{\partial U} \left[c_n K^{-1} \prod_{i=1}^n (K - z_i)^{-1} \right] \phi_2,
\end{aligned}$$

which can be simplified in a couple of ways.

We could express this force term as a function of the inverse $[KP(K)]^{-1}$ and then calculate $[KP(K)]^{-1} \phi$, but then we would require $n + 1$ matrix operations per conjugate gradient iteration, which negates the performance benefit gained from $KP(K)$ being close to unity.

A better solution is to expand the inverse polynomial into a sum over poles, then use a multi-shift solver to calculate the shifted inverses $[K - z_i]^{-1} \phi$. The general formula for this expansion is

$$\prod_{i=1}^n \frac{1}{K - z_i} = \sum_{i=1}^n \left(\prod_{j \neq i} \frac{1}{z_j - z_i} \right) \frac{1}{K - z_i} \equiv \sum_{i=1}^n \frac{r_i}{K - z_i}. \tag{C.6}$$

Taking $z_{n+1} = 0$, the force term becomes

$$\begin{aligned}
F_2 &= \phi_2^\dagger \frac{\partial}{\partial U} \left[\sum_{i=1}^{n+1} \frac{r_i}{K - z_i} \right] \phi_2 \\
&= - \sum_{i=1}^{n+1} \phi_2^\dagger [K - z_i]^{-1} \frac{\partial K}{\partial U} r_i [K - z_i]^{-1} \phi_2 \\
&= - \sum_{i=1}^{n+1} \bar{\eta}_i^\dagger \frac{\partial K}{\partial U} \bar{\chi}_i,
\end{aligned} \tag{C.7}$$

where

$$\bar{\chi}_i = r_i [K - z_i]^{-1} \phi_2, \quad (\text{C.8a})$$

$$\bar{\eta}_i = [K - z_i^*]^{-1} \phi_2, \quad (\text{C.8b})$$

and

$$r_i = \prod_{j=1, j \neq i}^{n+1} \frac{1}{z_j - z_i}. \quad (\text{C.9})$$

As $P(K)$ must be a real polynomial to avoid the sign problem, the roots z_i are either real or come in complex-conjugate pairs. Hence, for each z_i there is some z_j such that $z_i^* = z_j$, so we only need to construct $n + 1$ shifted inverses.

Appendix C.3. Mass preconditioning

Consider a mass preconditioned system

$$S_{\text{MP}} = \phi_1^\dagger J^{-1} \phi_1 + \phi_2^\dagger J K^{-1} \phi_2. \quad (\text{C.10})$$

The force term for the heaviest fermion F_1 is identical to that for basic HMC, namely

$$F_1 = -\phi_1^\dagger J^{-1} \frac{\partial J}{\partial U} J^{-1} \phi_1. \quad (\text{C.11})$$

As for the correction term, we have

$$\begin{aligned} F_2 &= \frac{\partial}{\partial U} [\phi_2^\dagger J K^{-1} \phi_2] \\ &= \phi_2^\dagger \frac{\partial J}{\partial U} K^{-1} \phi_2 - \phi_2^\dagger J K^{-1} \frac{\partial K}{\partial U} K^{-1} \phi_2. \end{aligned}$$

To make this look more symmetric, we expand $K = M^\dagger M$ and $J = W^\dagger W$ and write

$$F_2 = \phi_2^\dagger M^{-1} \frac{\partial J}{\partial U} (M^\dagger)^{-1} \phi_2 - \phi_2^\dagger W^\dagger K^{-1} \frac{\partial K}{\partial U} K^{-1} W \phi_2 \quad (\text{C.12})$$

Appendix C.4. Multiple filters

For actions with multiple polynomial and/or mass filters, each force term takes the form of one of the previously mentioned force terms: (C.4), (C.7), (C.11) or (C.12).

For example, the force terms for the PF-MP action

$$S_{\text{PF-MP}} = \phi_1^\dagger P(J) \phi_1 + \phi_2^\dagger [P(J)J]^{-1} \phi_2 + \phi_3^\dagger J K^{-1} \phi_3$$

are

$$F_1 = \phi_1^\dagger \sum_{i=1}^n \left[c_n \prod_{j=1}^{i-1} (J - z_j) \frac{\partial J}{\partial U} \prod_{j=i+1}^n (J - z_j) \right] \phi_1, \quad \sim (\text{C.4})$$

$$F_2 = - \sum_{i=1}^{n+1} \phi_2^\dagger [J - z_i]^{-1} \frac{\partial J}{\partial U} r_i [J - z_i]^{-1} \phi_2, \quad \sim (\text{C.7})$$

and

$$F_3 = \phi_3^\dagger M^{-1} \frac{\partial J}{\partial U} (M^\dagger)^{-1} \phi_3 - \phi_3^\dagger W^\dagger K^{-1} \frac{\partial K}{\partial U} K^{-1} W \phi_3. \sim (\text{C.12})$$

References

References

- [1] S. Duane, A. D. Kennedy, B. J. Pendleton, D. Roweth, Hybrid Monte Carlo, Phys. Lett. B195 (1987) 216–222. [doi:10.1016/0370-2693\(87\)91197-X](#).
- [2] A. Ukawa, Computational cost of full QCD simulations experienced by CP-PACS and JLQCD Collaborations, Nucl. Phys. Proc. Suppl. 106 (2002) 195–196. [doi:10.1016/S0920-5632\(01\)01662-0](#).
- [3] M. Hasenbusch, Speeding up the hybrid Monte Carlo algorithm for dynamical fermions, Phys. Lett. B519 (2001) 177–182. [arXiv:hep-lat/0107019](#), [doi:10.1016/S0370-2693\(01\)01102-9](#).
- [4] W. Kamleh, M. Peardon, Polynomial Filtered HMC: An Algorithm for lattice QCD with dynamical quarks, Comput. Phys. Commun. 183 (2012) 1993–2000. [arXiv:1106.5625](#), [doi:10.1016/j.cpc.2012.05.002](#).
- [5] M. Lüscher, Solution of the Dirac equation in lattice QCD using a domain decomposition method, Comput. Phys. Commun. 156 (2004) 209–220. [arXiv:hep-lat/0310048](#), [doi:10.1016/S0010-4655\(03\)00486-7](#).
- [6] M. Clark, A. Kennedy, Accelerating dynamical fermion computations using the rational hybrid Monte Carlo (RHMC) algorithm with multiple pseudofermion fields, Phys. Rev. Lett. 98 (2007) 051601. [arXiv:hep-lat/0608015](#), [doi:10.1103/PhysRevLett.98.051601](#).
- [7] J. C. Sexton, D. H. Weingarten, Hamiltonian evolution for the hybrid Monte Carlo algorithm, Nucl. Phys. B380 (1992) 665–677. [doi:10.1016/0550-3213\(92\)90263-B](#).
- [8] I. Omelyan, I. Mryglod, R. Folk, Symplectic analytically integrable decomposition algorithms: classification, derivation, and application to molecular dynamics, quantum and celestial mechanics simulations, Comput. Phys. Commun. 151 (2003) 272–314. [doi:10.1016/S0010-4655\(02\)00754-3](#).
- [9] C. Urbach, K. Jansen, A. Shindler, U. Wenger, HMC algorithm with multiple time scale integration and mass preconditioning, Comput. Phys. Commun. 174 (2006) 87–98. [arXiv:hep-lat/0506011](#), [doi:10.1016/j.cpc.2005.08.006](#).
- [10] S. Aoki, et al., Physical Point Simulation in 2+1 Flavor Lattice QCD, Phys. Rev. D81 (2010) 074503. [arXiv:0911.2561](#), [doi:10.1103/PhysRevD.81.074503](#).
- [11] M. Bruno, et al., Simulation of QCD with $N_f = 2 + 1$ flavors of non-perturbatively improved Wilson fermions, JHEP 02 (2015) 043. [arXiv:1411.3982](#), [doi:10.1007/JHEP02\(2015\)043](#).
- [12] R. Arthur, et al., Domain Wall QCD with Near-Physical Pions, Phys. Rev. D87 (2013) 094514. [arXiv:1208.4412](#), [doi:10.1103/PhysRevD.87.094514](#).
- [13] M. J. Peardon, J. Sexton, Multiple molecular dynamics time scales in hybrid Monte Carlo fermion simulations, Nucl. Phys. Proc. Suppl. 119 (2003) 985–987. [arXiv:hep-lat/0209037](#), [doi:10.1016/S0920-5632\(03\)01738-9](#).
- [14] A. Ali Khan, et al., Accelerating the hybrid Monte Carlo algorithm, Phys. Lett. B564 (2003) 235–240. [arXiv:hep-lat/0303026](#), [doi:10.1016/S0370-2693\(03\)00703-2](#).

- [15] S. Borsanyi, S. Durr, Z. Fodor, C. Hoelbling, S. Katz, et al., Ab initio calculation of the neutron-proton mass difference, *Science* 347 (2015) 1452–1455. [arXiv:1406.4088](#), [doi:10.1126/science.1257050](#).
- [16] M. A. Clark, B. Joó, A. D. Kennedy, P. J. Silva, Improving dynamical lattice QCD simulations through integrator tuning using Poisson brackets and a force-gradient integrator, *Phys. Rev. D* 84 (2011) 071502. [arXiv:1108.1828](#), [doi:10.1103/PhysRevD.84.071502](#).
- [17] Y. Nakamura, H. Stüben, BQCD - Berlin quantum chromodynamics program, *PoS LATTICE2010* (2010) 040. [arXiv:1011.0199](#).
- [18] Lüscher, Martin, Deflation acceleration of lattice QCD simulations, *JHEP* 12 (2007) 011. [arXiv:0710.5417](#), [doi:10.1088/1126-6708/2007/12/011](#).
- [19] T. Takaishi, Choice of integrator in the hybrid Monte Carlo algorithm, *Comput. Phys. Commun.* 133 (2000) 6–17. [arXiv:hep-lat/9909134](#), [doi:10.1016/S0010-4655\(00\)00161-2](#).
- [20] C. Gattringer, C. Lang, *Quantum Chromodynamics on the Lattice: An Introductory Presentation*, Vol. 788 of *Lecture Notes in Physics*, Springer Berlin Heidelberg, 2010. [doi:10.1007/978-3-642-01850-3](#).
- [21] J. Osborn, [The FUEL code project](#), *PoS LATTICE2014* (2014) 028.
URL https://pos.sissa.it/archive/conferences/214/028/LATTICE2014_028.pdf